School Choice and Information
An Experimental Study on Matching
Mechanisms

Joana Pais[*], Ágnes Pintér[†]

May, 2006

Abstract

We present an experimental study where we analyze three well-known matching mechanisms—the Boston, the Gale-Shapley, and the Top Trading Cycles mechanisms—in three different informational settings. Our experimental results are consistent with the theory, suggesting that the TTC mechanism outperforms both the Boston and the Gale-Shapley mechanisms in terms of efficiency and it is as successful as the Gale-Shapley mechanism regarding the proportion of truthful preference revelation, whereas manipulation is stronger under the Boston mechanism. In addition, even though agents are much more likely to revert to truthtelling in lack of information about the others' payoffs—ignorance may be beneficial in this context—, the TTC mechanism results less sensitive to the amount of information that participants hold. These results therefore suggest that the use of the TTC mechanism in practice is more desirable than of the others.

1   Introduction

There is now a vast literature on matching problems. Matching is a pervasive phenomenon arising in several economic and social settings. The assignment of civil servants to civil service positions, the admission of students to colleges, some entry-level labor markets—as the widely explored market for graduating physicians—, or the school choice problem are among the matching situations that have gained attention in the last decades. The working of some matching mechanisms, along with strategic issues that confront individuals in these contexts, have been explored theoretically.

Very briefly, in a two-sided matching market, agents belong to one of two disjoint sets, say colleges and students, and each agent—college and student—has preferences over the other side of the market—students and colleges, respectively—and the prospect of being unmatched. The matching problem then reduces to assigning students to colleges by means of a matching mechanism. Stability, strategy-proofness, and Pareto efficiency of such mechanisms are highly valued properties. A mechanism is stable if it always selects stable matchings; by definition, under a stable matching every agent in the market prefers his partner to being alone and, moreover, no pair of agents—consisting of a college and a student—who are not matched to each other would rather prefer to be so matched. A mechanism is strategy-proof if it is immune to preference manipulation, i.e., truth is a dominant strategy. A mechanism is Pareto efficient if it always selects Pareto efficient matchings.

The perhaps most famous matching mechanism relies on the Gale-Shapley deferred-acceptance algorithm (Gale and Shapley, 1962). Gale and Shapley were motivated by the problem of the admission of students to colleges and the Gale-Shapley algorithm was written as a means to show that a stable matching always exists in such a two-sided matching market. The Gale-Shapley deferred-acceptance algorithm transforms a matching where all agents are unmatched into a stable matching, thus proving existence. Besides guaranteeing stability, the Gale-Shapley mechanism has other appealing properties. Namely, truth is a dominant strategy for one side of the market (Dubins and Freedman, 1981, Roth, 1982a). Moreover, it is Pareto efficient when the welfare of both sides of the market is considered (Roth, 1982a).

The strategic properties of the Gale-Shapley mechanism—namely the characterization of dominant strategies and equilibria—and, in fact, most theoretical studies on matching mechanisms rely on the assumption of complete information, however implausible: knowing the true preferences of every agent in the market is more than we may reasonably expect in most matching markets. Only a few papers have relaxed this assumption and are thus worth mentioning. Roth (1989) is a first attempt to deal with the incomplete

information case. Under incomplete information, even though truth remains a dominant strategy for one side of the market when the Gale-Shapley mechanism is employed, the equilibrium characterization for the complete information case is not robust. Ehlers and Massó (2003) study Bayesian Nash equilibria for mechanisms producing stable matchings—as the Gale-Shapley mechanism—and find a necessary and sufficient condition for truthtelling to be an equilibrium: truthtelling is a Bayesian Nash equilibrium in the revelation game induced by a stable mechanism and a common belief if and only if all profiles in its support have a singleton core. Finally, Roth and Rothblum (1999) and Ehlers (2003, 2004) are less ambitious and do not aim at characterizing equilibria, but give advice to individuals on how to participate in matching markets when there is uncertainty about the others' strategies.

Still, many questions regarding the strategic incentives agents face under incomplete information remain to be answered on theoretical grounds. How the amount of information held by individuals on the elements of the game actually influences individuals' decision making, affecting the performance of matching mechanisms, is thus a question to be explored. For instance, Barberà and Dutta (1995) consider truthtelling as a form of "protective" behavior, claiming that risk averse agents may revert to faithfully revealing their true preferences when they are poorly informed. Moreover, it is clear that in mechanisms for which truth is not a dominant strategy, computing the optimal strategies requires a lot of information on others' preferences. In this paper we present an experimental study to investigate these and other issues, providing a direction into which the role of information on decision making may be ascertained.

We investigate a particular class of matching problems: the assignment of individuals to indivisible items. In these problems, individuals—let us call them teachers—have strict preferences over the indivisible items—henceforth, schools—and, on the other hand, schools have a maximum capacity and a strict priority ordering of all teachers. This problem has been referred to as the school choice problem (Abdulkadiroglu and Sönmez, 2003) and is closely related to the college admissions problem explored by Gale and Shapley (Gale and Shapley, 1962), the main difference being that, in contrast to the college admissions model, here schools are not strategic agents, but mere objects to be assigned to teachers. Hence, while teachers may not straightforwardly reveal their true preferences, schools have no chance of manipulating priorities.

The influence of information is assessed for the Gale-Shapley mechanism and for another well-known matching mechanism, the Top Trading Cycles mechanism (TTC), as well as for the Boston mechanism, which has been

widely used in the context of school choice problems.[1] The TTC (Shapley and Scarf, 1974) fulfills two appealing properties: it is both strategy-proof (Roth, 1982b) and Pareto efficient. The Gale-Shapley mechanism is strategy-proof, but not efficient (Roth, 1982a), since we only consider teachers' welfare in this setup. Finally, we have included the Boston mechanism for reference, as it fails to meet both requirements: it is neither strategy-proof nor Pareto efficient. In fact, there is room for profitable manipulation of agents' preferences under this mechanism and misrepresentation leads to major efficiency losses.[2]

Besides providing yet another test of theoretical results on matching mechanisms with boundedly rational individuals, we address two main questions. First, we compare the three above mentioned mechanisms under three informational scenarios, ranging from complete ignorance about the other participants' preferences and schools' priorities to complete information on all elements of the game. In particular, we are concerned in comparing the incentives agents face under different mechanisms, as well as in comparing efficiency levels, for different information levels. The results in this paper seem to suggest that the TTC mechanism prevails over the other two in what efficiency is concerned, even though incentives for truthtelling are not significantly different from those given by Gale-Shapley. Second, within each mechanism, we evaluate the influence of the amount of information held by individuals on decision making. Namely, we are concerned in testing whether truthtelling emerges as a very salient form of behavior when information is low. This would imply that, in markets where we expect agents to know little about others, strategy-proofness of the mechanism should not drive the choice of the matching mechanism. Moreover, if we are able to determine that information significantly affects individuals' behavior, we can immediately conclude that the existing theoretical results, which rely on the fundamental assumption of complete information, are insufficient to deal with markets where agents know little about others. Otherwise, if the effect of information is not relevant, theory may be considered apt to deal with the incomplete information case. Our results support the first conjecture: information is actually important. In fact, in a very low information environment, acting straightforwardly is a very salient form of behavior and significantly higher levels of efficiency are achieved under all mechanisms except for TTC, which

---

[1]School choice programs have become increasingly popular in the US. The best known of these programs rely on the Boston mechanism, used to assign students to schools in Boston, Cambridge, and Seattle, among others.

[2]On the functioning and strategic properties of the Boston mechanism, check Abdulkadiroglu and Sönmez (2003), Abdulkadiroglu et al. (2005), Chen and Sönmez (2004), and Ergin and Sönmez (2004).

appears to be less sensitive to the amount of information held by participants.

We are aware of several experimental studies of matching problems, some of which aim at testing the above mentioned mechanisms. These include Harrison and McCabe (1996) that explores the Gale-Shapley mechanism and shows that profitable manipulation of agents' preferences becomes more difficult as markets get larger; Chen and Sönmez (2002a) that compares a random serial dictatorship mechanism used to allocate dormitory rooms in American universities with a variant of the TTC in an incomplete information environment, concluding that the TTC produces significantly more efficient allocations; in a companion paper, Chen and Sönmez (2002b) evaluate the performance of these mechanisms under complete information, reaching the same qualitative results; finally, Chen and Sönmez (2004), consider the school choice problem and analyze the TTC, the Gale-Shapley, and the Boston mechanisms under incomplete information, concluding that, in what efficiency is concerned, Gale-Shapley improves upon the TTC, which outperforms the Boston mechanism. The difference between the above studies and this paper derives from our main objective: to test the role of information in evaluating matching mechanisms. Other experimental studies, dealing with other matching mechanisms, are: Olson and Porter (1994), Nalbantian and Schotter (1995), Kagel and Roth (2000), Ünver (2001), Haruvy, Roth, and Ünver (2001), and McKinney, Niederle, and Roth (2005).

We proceed as follows. In Section 2 we present the theoretical properties of the three matching mechanisms under study. We describe the experimental design in Section 3. Section 4 summarizes the main results of the experiments. Some concluding remarks follow in Section 5.

2   The Theoretical Model

We first introduce the model and then describe the three matching mechanisms and their theoretical properties.

In this assignment problem there are a number of teachers to fill a number of vacancies or teaching positions across different schools. Each teacher has strict preferences over all schools, while each school has a strict priority ranking of all teachers, as well as a maximum number of teachers to employ. Priorities are exogenous and not subject to manipulation by schools. The fact that only teachers can act strategically is what distinguishes this problem from the college admissions model. It is not difficult to justify the use of the school choice model. Besides the fact that this model is easier to be implemented in the laboratory, we can find plenty of real-life situations that can be described as one side of the market being inactive. We have already

mentioned the use of the school choice problem in the admission of children to public schools in the US, which also applies to other countries as Spain, but we can also think about the admission of students to universities (which is, in most countries, based on students' grades), the MIR system ("Médicos Internos Residentes," a residence training system for physicians in public hospitals based on their performance) in Spain, or in general the assignment of civil servants to civil service positions, which is, in several countries, based on an objective scoring system (for example teachers, judges, or tax inspectors in Spain).

The outcome of the school choice problem is a matching, an assignment of teachers to teaching positions such that each teacher is assigned one vacancy and each vacancy is filled by one teacher only. A matching is Pareto efficient if there is no matching that assigns at least one teacher a strictly better school and every other teacher a weakly better school. A matching mechanism consists of a systematic procedure that selects a matching for each school choice problem. A matching mechanism is efficient if it always chooses Pareto efficient matchings; it is strategy-proof if truth is a dominant strategy, i.e., no teacher can profitably manipulate her preferences, independently of the other agents' strategies.

2.1 The Top Trading Cycles Mechanism

In this context, the TTC works as follows:

1. Each school gives priority to a number of teachers up to its capacity; in this setting, for simplicity, each teacher has priority in one school only.

2. Each teacher reports her preferences over the schools.

3. An ordering of teachers is randomly chosen.

4. For any submitted teachers' preferences, schools' priorities, and ordering of teachers, the outcome is obtained after undergoing the following steps:

   (a) Assign each teacher to a school (tentative assignment); in this setting, each teacher is tentatively assigned to her priority school.

   (b) The first teacher in the ordering proposes to her top ranked school. If she has priority at this school, the assignment is finalized and both the teacher and teaching position are removed from the system; the procedure continues with the second teacher in the ordering. Otherwise, the first teacher in the ordering that is tentatively

assigned to the proposed school is inserted at the top of the ordering, in front of the requester.

(c) When the ordering is modified, this procedure is repeated, so that the teacher who just became first in the ordering sends an application to her highest-ranked school. If she has priority at this school, the assignment is finalized and the procedure continues with the next teacher in line. Otherwise, the first teacher in the ordering tentatively assigned to the proposed school is inserted at the top of the ordering, in front of the requester.

(d) If a cycle forms, it consists of a sequence of proposals of the kind: A proposes to B's tentative assignment, B applies to C's tentative assignment, and C proposes to A's tentative assignment. In such cases, all teachers in the cycle are assigned to the schools they proposed to and teachers, as well as their respective assignments, are removed from the system.

(e) The procedure stops when all teachers are assigned to a position.

The TTC mechanism satisfies two appealing properties: it is strategy-proof, i.e., truth is a dominant strategy for every teacher, and Pareto efficient.[3] We thus expect that individuals reveal their preferences in a straightforward manner, independently of the amount of information they hold on the elements of the game.[4]

2.2 The Gale-Shapley Mechanism

The Gale-Shapley is certainly one of the best known mechanisms in the matching literature.

Its theoretical properties and the incentives it gives to agents have been scrutinized and its applications encompass a significant number of markets. In what follows we describe the functioning of the Gale-Shapley mechanism:

1. A priority ordering of teachers is determined for each school.

2. Each teacher reports her preferences over the schools.

3. Given the submitted preferences of the teachers and schools' priority orderings, positions are allocated after undergoing the following steps:

---

[3]Note also that, in this setting, the resulting final assignment is independent of the random ordering of teachers defined in step 3.

[4]Still, Chen and Sönmez (2002a) find, in an experiment about on-campus housing, that about one-third of the subjects manipulate their preferences under a variant of the TTC mechanism.

(a) Each teacher proposes to her first ranked school. Each school keeps the applicants with higher priority order on hold until positions are filled, while rejecting the lowest priority teachers in excess of its capacity.

(b) In general:

Every teacher who got rejected in the previous step proposes to the next school on her list of preferences. Each school considers the teachers it holds from the previous step together with the new applications. The lowest priority teachers in excess of the school's capacity are rejected, while remaining applications are kept on hold.

(c) This process is repeated until no applications are rejected. Each participant is then assigned the position at the school that keeps her on hold.

As the TTC mechanism, the Gale-Shapley mechanism is strategy-proof. Again, we expect individuals to faithfully reveal their true preferences over schools in every informational treatment. This mechanism is efficient when the welfare of both sides of the market is taken into account. Nevertheless, in this assignment problem, schools are mere objects to be allocated among teachers and only teachers' welfare is taken into consideration. Since there may exist a matching that Pareto dominates the outcome of the Gale-Shapley mechanism for teachers, the mechanism is not efficient in this setup. It follows that, if theory is to be confirmed, the TTC should outperform the Gale-Shapley in efficiency terms.

2.3 The Boston Mechanism

The Boston mechanism has been the most widely used assignment mechanism in real-life applications of school choice problems. It works as follows:

1. A priority ordering of teachers is determined for each school.

2. Each teacher reports her preferences over the schools.

3. Given the submitted preferences of the teachers and schools' priority orderings, positions are allocated after several rounds:

(a) Each teacher proposes to her top ranked school. Each school accepts the proposals from the teachers with higher priority order until positions are filled (or no teachers remain who have proposed

to the school). These applicants and their positions are removed from the system. All other applications are rejected by the schools.

(b) In general at round k:

Each teacher remaining in the system proposes to its k^{th} school. Each school with vacant positions accepts the proposals from the teachers with higher priority order until positions are filled (or no teachers remain who have proposed to the school in this round). These applicants and their positions are removed from the system. All other applications are rejected by the schools.

(c) The procedure terminates when each teacher is assigned a position.[5]

A major handicap of the Boston mechanism is that it leads to preference manipulation.[6] In fact, teachers are given incentives to rank high on their submitted preferences the schools where they have good chances of getting in. This has two important consequences. First, evaluating the performance of this mechanism according to the revealed preferences is clearly inadequate. Moreover, even though the outcome of the Boston mechanism is Pareto efficient when teachers submit their true preferences, preference manipulation may lead to a substantial efficiency loss. Hence, we expect a low level of efficiency.

3   Experimental Design

We design our experiments to analyze participants' decision taking under different informational settings throughout the above described matching mechanisms: the Boston mechanism, the Gale-Shapley mechanism, and the TTC mechanism. We use a 3x3 design: for each mechanism we construct three treatments differing in the amount of information held by participants about the elements of the game. This allows us to compare decision making in nine treatments. In our analysis we concentrate on the role of information in truthful preference revelation and in efficiency. The environment is designed to capture the key aspects and difficulties of each mechanism, under a controlled environment, with relatively small groups of participants.

Participants are randomly and anonymously sorted into groups of five. Each participant plays the role of a teacher to be assigned to a teaching

---

[5]Hence, if there are n teachers and l schools, the process ends in a maximum number of l rounds.

[6]Chen and Sönmez (2003) report an experiment where around 80% of the subjects manipulate their true preferences.

position. For each group of five teachers, there are five vacancies—or teaching positions—across three schools that differ in capacity (number of opening positions) and desirability. Each position should be assigned to one teacher only. Preferences over schools are induced by the monetary payoff a teacher obtains depending on the school where she fills a vacancy at the end of the experiment. The payoffs obtained are symmetric: every teacher gets 15€ for her top choice, 9€ for the second choice, and 3€ for the last choice, but different teachers need not agree on which school is either her top, second, or last choice. The payoffs of different outcomes are sufficiently dispersed so as to have a monetarily salient difference (12€) between getting one's best and one's worst choice.

New groups are formed for each mechanism by shuffling all participants under the "role" of a particular teacher. Participants are informed that new groups are created for each mechanism. Under each mechanism, each group of participants plays the game three times, holding three different, but increasing, amounts of information. As a means to avoid participants from playing from memory when the mechanism is changed, we modify the labeling of schools and alter the schools' characteristics accordingly. It follows that each participant keeps the same preferences over schools throughout the whole experiment, even though her payoff matrix seemingly changes: the performed changes are virtual.

Finally, schools have priorities over teachers. This means that schools may prefer some applicants to others and are able to rank all the participants in a list of priorities. Moreover, as priorities of schools are given, schools are not real strategic agents (i.e., they "play" truthfully) and all the participants know this.

3.1   Informational Settings

In each experimental session three different informational treatments are implemented for each mechanism, in the following order:

- **Zero information setting**: In this setting each participant knows her possible payoff amounts depending on the school where she holds a position (i.e., her own induced preferences), but not the other participants' preferences. She is only told that different participants might have different payoff tables.

  Participants have no information about the schools' priority ordering in this treatment. They are only told the capacity (i.e., the number of vacancies) of each school.

10

- **Partial information setting**: In this setting, besides her own induced preferences and the capacity of each school, each participant has some partial information about the schools' priority orderings. Namely, each participant is told the favorite candidates of each school, up to its capacity.

- **Full information setting**: In this informational setting each participant has complete information on both the induced preferences of all participants, and the full priority ordering of schools over candidates.

In the case of the TTC mechanism, as the schools' priority orderings are reflected in the tentative assignment, under the "partial information treatment" participants are told—besides their own induced preferences—the tentative assignment of all participants; while in the "full information treatment" they know both the induced preferences of all participants and the tentative assignment.

We conducted sessions with both undergraduate and graduate students from the Universitat Autònoma de Barcelona, recruited using classroom announcements and posters in the campus, where the experimental sessions (on paper/by hand) took place. In total, 30 subjects have participated in the experiment, which makes six groups of five. Each subject was allowed to participate in one session only. Subjects were informed that they would participate in a decision making task. At the beginning of each session, subjects were randomly seated at the tables and printed instructions were given to them. Before starting each mechanism, the corresponding instructions were read aloud. In each session and for each informational setting participants were asked to submit a school ordering, from their top to their last choice. In each session all nine treatments were implemented, so that we have nine decisions from each participant, each under a different mechanism and/or different informational setting. The order of implementation of the mechanisms and informational settings was the following: 1^{st}: Boston mechanism in the zero information setting, 2^{nd}: Boston mechanism in the partial information setting, 3^{rd}: Boston mechanism in the full information setting; 4^{th}: Gale-Shapley mechanism in the zero information setting, 5^{th}: Gale-Shapley mechanism in the partial information setting, 6^{th}: Gale-Shapley mechanism in the full information setting, 7^{th}: TTC mechanism in the zero information setting, 8^{th}: TTC mechanism in the partial information setting, and finally 9^{th}: TTC mechanism in the full information setting. Subjects did not get any feedback about previous decisions or outcomes at any moment of the experiment. At the end of each session, one of the nine treatments was chosen randomly for payment, matching was determined for the chosen treatment

and earnings were paid. Sessions lasted about 90 minutes and average net payments—including a 3€ show-up fee—were around 14€.

The instructions and Decision Sheets in English for the Gale-Shapley mechanism can be found in the Appendix.[7]

4   Experimental Results

In this section we present our experimental results. Our main aim is to examine the importance of the level of information that participants hold in evaluating the three matching mechanisms. The experimental setting we use allows us to analyze the role of information in individuals' decision making and how this affects the theoretic characteristics of the mechanisms. The first keypoint is related to whether individuals report their preferences truthfully. We study whether the amount of information given to the participants influences truthful preference revelation (keeping the mechanism under analysis fixed) and, on the other hand, whether under the same informational setting truthful preference revelation changes with the implemented mechanism. The second main point is related to efficiency. We compare efficiency levels under each mechanism across information settings and across different mechanisms for the same information level.

4.1   Truthful Preference Revelation

We first analyze the two questions related to truthful preference revelation. Table 1 shows the proportion of participants who played truthfully (regarding induced preferences) and who used three possible kinds of preference manipulation, in each treatment.

---

[7]The instructions for the other mechanisms only differ in the description of the allocation method. The Decision Sheets for players in different roles look similar to the ones shown in the Appendix.

Mechanism		Info	Not.	Truth	PSB	SSB	PSB&SSB	Other
B	**o**	Zero	B0	86.7%	-	10%	0%	3.3%
s	**t**	Partial	B1	33.3%	36.7%	6.7%	16.7%	6.7%
o	**n**	Full	B2	36.7%	43.3%	3.3%	3.3%	13.3%
G a l e	Sh ap le y	Zero	GS0	80%	-	6.7%	0%	13.3%
l e	p l	Partial	GS1	60%	30%	0%	3.3%	6.7%
&	e y	Full	GS2	66.7%	10.0%	6.7%	6.7%	10%
T		Zero	T0	96.7%	-	3.3%	0%	0%
T		Partial	T1	63.3%	23.3%	10%	3.3%	0%
C		Full	T2	63.3%	16.7%	10%	0%	10%

Table 1: Proportion of truthful play and preference manipulations

We can see that the proportion of players who played truthfully varies between 33,3% and 96,7%, depending on the treatment being implemented. Even though the Gale-Shapley and the TTC mechanisms are strategy-proof, there is some misrepresentation of preferences in these treatments as well, although its extent depends on the implemented information setting. It is important to examine who manipulates the preferences and in which manner. We identify three possible ways of preference manipulations. First, a substantial proportion of the participants has ranked the school where she has priority higher in the submitted ranking than it would be according to the induced preferences; this is what we call the "Priority School Bias" (PSB).[8] The second identified way of manipulating the true preferences is to underrank the most competitive school (i.e., the school with only one vacancy); following Chen and Sönmez (2004) we call this form of behavior "Small School Bias" (SSB). The third manipulation method (PSB&SSB) is simply the simultaneous use of both previously described ways.[9] We can see that a relatively small proportion of the participants used any of these two latter methods, and the most frequently used method is the PSB. Checking the average payoff of the participants that use this method (8.23 euro), we can

---

[8]In the case of the zero information setting it does not make sense to check for this kind of preference manipulation, as participants in these settings have no information about the priorities of the schools.

[9]As we have already mentioned, in the case of the zero information setting it does not make sense to check for the PSB preference manipulation, therefore in these info settings the PSB&SSB manipulation will be included in the SSB manipulation method.

infer that PSB is used as a means to ensure their second best payoff (9 euro). This form of behavior can be considered as a kind of risk aversion, since this manipulation allows the participants to avoid their worst payoff (3 euro). Although the SSB and the PSB&SSB methods are used less frequently, they yield a slightly higher payoff to the participants than the previous method.

Now, we would like to analyze whether the quantity of information the participants hold affects behavior. To answer to this question, we compare the proportion of participants playing truthfully in different informational settings under each mechanism.[10]

Result 1: *Under each examined matching mechanism (Boston, Gale-Shapley, and TTC) the amount of information has a significant effect on the decisions of the participants in what truthtelling is concerned. In particular, we find that having no information about the other parties' preferences results in a significantly higher proportion of truthtelling than under any treatment with additional information.*

Statistical evidence: Under each mechanism, the null hypothesis of equal proportions of truthful preference revelation across the three informational settings can be rejected at 5% significance level. In particular, in the Boston and TTC mechanisms, the null can be rejected at any reasonable significance level and in Gale-Shapley at 3.7%.[11] As the null hypotheses are rejected, multiple comparisons are made. The results of the pairwise comparisons (with the corresponding test statistics of the Quade test[12] in parenthesis) can be found in Table 2:

Mechanism	Boston			Gale&Shapley			TTC		
Null hyp.	B0=B1	B0=B2	B1=B2	GS0=GS1	GS0=GS2	GS1=GS2	T0=T1	T0=T2	T1=T2
Difference	53.3%*	50%*	3.3%	20%*	13.3%**	7%	33.3%*	33.3%*	0%
	(28.8)	(34.3)	(5.5)	(27)	(18)	(9)	(30.5)	(32.5)	(2)

* Significant at α=0.01   ** Significant at α=0.1

Table 2: Differences in the proportion of truthtelling across information settings

---

[10]We compare, for each mechanism, three informational settings: B0 vs B1 vs B2; GS0 vs GS1 vs GS2; and T0 vs T1 vs T2.

[11]The value of the test statistics of the Quade test for k related samples in the Boston mechanism is 8.661, in Gale-Shapley it is 5.385, and in TTC it is 11.479; while the 0.01 quantile of the F distribution with k_1=2 and k_2=10 degrees of freedom is 7.56, and the 0.037 quantile is 5.26.

[12]To learn more about the Quade test used in our analysis, see W.J. Conover (1980).

From the results of the comparisons we can conclude that under the Boston and TTC mechanisms the proportion of truthful preference revelation is significantly higher at 1% significance level in the zero information setting than either in the partial or the full information settings (86.7% vs. 33.3% and 36.7% respectively under the Boston mechanism; 96.7% vs. 63.3% and 63.3% respectively under the TTC mechanism), while there is no significant difference in the proportions between the partial and full information settings. Under the Gale-Shapley mechanism the proportion of truthful preference revelation is significantly higher at 1% significance level in the zero information setting than in the partial information setting (80% vs. 60%), but its difference with the full information setting is only significant at 10% (80% vs. 66.7%). Here again there is no significant difference in the proportion of straightforward behavior between the partial and the full information settings.

Result 2: *Across mechanisms we find a significant difference in the proportion of truthtelling. In particular, both in the partial and full information settings, the proportion of truthtelling under the Boston mechanism is lower than under any of the other mechanisms. In the case of the zero information setting, under the TTC mechanism, the proportion of truthful preference revelation is significantly higher than under any other mechanism.*

Statistical evidence: Under the three informational settings, the null hypothesis of equal proportions of truthful preference revelation across the three matching mechanisms can be rejected at 5% significance level. In particular, in both the zero and full information settings the null can be rejected at any reasonable significance level, while in the partial information setting it can be rejected at 3.4%.[13] As the null hypotheses are rejected, multiple comparisons are made. The results of the pairwise comparisons can be found in Table 3:

---

[13]The value of the test statistics of the Quade test for k related samples in the zero information setting is 8.636, in the partial information setting it is 5.526, and in the full information setting it is 7.986; while the 0.01 quantile of the F distribution with k_1=2 and k_2=10 degrees of freedom is 7.56, and the 0.034 quantile is 5.484.

Info	Zero			Partial			Full		
Null hyp.	B0=GS0	T0=B0	T0=GS0	GS1=B1	T1=B1	T1=GS1	GS2=B2	T2=B2	GS2=T2
Difference	7%	10%**	16.7%*	26.7%**	30%**	3%	30%*	26.7%*	3%
	(12)	(18)	(30)	(24)	(30)	(6)	(31.3)	(27.3)	(4)

* Significant at α=0.01   ** Significant at α=0.05

Table 3: Pairwise differences in truthtelling across mechanisms

A significantly smaller number of participants reveal their preferences truthfully in the partial and the full information settings under the Boston mechanism (at 5% significance level) than under either the Gale-Shapley or TTC mechanism (in the partial information setting 33% vs. 60% and 63% respectively; and in the full information setting 37% vs. 67% and 63% respectively), while there is no significant difference between the Gale-Shapley and TTC mechanisms.[14] As for the case of the zero information setting, our results show that under the TTC mechanism the proportion of truthful preference revelation is significantly higher (at 5% significance level) than under either the Gale-Shapley or the Boston mechanism (97% vs. 80% and 87%),[15] while there is no significant difference in truthtelling under the Boston and the Gale-Shapley mechanisms.

As we have seen above, participants manipulate their preferences in order to increase the chances of getting their top ranked choices, which do not necessarily coincide with their true top choices. Hence, the more participants get their reported but not true top choices, the lower the achieved efficiency. Therefore, it is worth to examine whether under some mechanisms and/or information settings the participants can get their true top choices.

Result 3: *Under each mechanism, in both the partial and full information settings, there is a highly significant difference between the proportion of the participants who receive their reported top choice and those who receive their true top choices. The proportion of subjects who receive their true top choices is significantly higher in the zero information setting than in any other information setting. On the other hand, across mechanisms the TTC mechanism results more successful in terms of the proportion of participants who get their true top choices.*

---

[14]In the case of the full information setting the same result holds already at 1% significance level.

[15]The difference of proportions between the TTC and Gale-Shapley mechanisms is significant already at the 1% significance level.

Statistical evidence: In the zero information setting, under all mechanisms, there is no significant difference between the proportion of subjects who get their reported top choices and the ones who get their true top choice. Using a t-test of proportions, the null hypothesis of equal proportions can not be rejected at any reasonable significance level, in favor of the alternative hypothesis of the percentage of reported top choices being higher than the percentage of true top choices.[16] In the case of the partial and full information settings, the same null hypothesis can be rejected at any reasonable significance level; therefore in these cases we can conclude that a significantly higher proportion of participants receive their reported top choices rather than their true top choices.[17]

On the other hand, under each mechanism there is no significant difference in the proportion of subjects that receive their reported top choices in the zero, partial, or full informational settings (under the Boston mechanism: B0: 83.3% vs. B1: 76.7% vs. B2: 80%; under Gale-Shapley: GS0: 66.7% vs. GS1: 50% vs. GS2: 66.7%; and under TTC: TTC0: 80% vs. TTC1: 80% vs. TTC2: 83.3%). If we consider the proportion of participants who receive their true top choices, using again a t-test of proportions we arrive to the conclusion that the null of equal proportions across information settings can be rejected. In particular, we find that under each mechanism a significantly higher proportion of participants receive their true top choices in the zero information setting, than in either the partial or the full information setting, at 10% significance level. Examining the proportion of participants who get their true top choices across mechanisms, we arrive to the conclusion that both in the partial and in the full information settings, under the TTC mechanism, a significantly higher proportion of the participants get their true top choices than either under the Boston or under the Gale-Shapley mechanism (at 10% significance level). In the zero information setting, the Boston and the TTC mechanisms result significantly more successful than the Gale-Shapley mechanism, in the same aspect. The following table contains

---

[16]The value of the test statistics in the Boston mechanism is 0.647, in Gale-Shapley it is 1.643, and in TTC it is 0.216; while the 0.01 quantile of the standard normal distribution is 2.327.

[17]The relevant proportions are the following:

Mechanism	Boston		Gale&Shapley		TTC	
Info setting	B1	B2	GS1	GS2	T1	T2
Reported top choice %	76.7%	80.0%	50.0%	66.7%	80.0%	83.3%
True top choice %	23.3%	40.0%	20.0%	36.7%	50.0%	53.3%

Each difference between the reported and true top choices in a treatment is significant, with a p-value <0.001.

the differences between information settings and across mechanisms, and indicates whether a difference is statistically significant.[18]

Difference	B0-B1	B0-B2	B2-B1	GS0-GS1	GS0-GS2	GS2-GS1	T0-T1	T0-T2	T2-T1
Diff. true top choice	53.3%*	36.7%*	16.7%**	30%*	13.3%***	16.7%**	26.7%*	23.3%*	3.3%
(across info settings)	(6.69)	(3.91)	(1.94)	(3.88)	(1.33)	(2.05)	(2.74)	(2.38)	(0.18)
Difference	B0-GS0	B0-T0	T0-GS0	B1-GS1	T1-B1	T1-GS1	B2-GS2	T2-B2	T2-GS2
Diff. true top choice	26.7%*	0%	26.7%*	3.3%	26.7%*	30%*	3.4%	13.3%***	16.7%**
(across mechanisms)	(2.74)	(-0.22)	(2.74)	(0.23)	(3.24)	(3.88)	(0.19)	(1.3)	(1.7)

* Difference significant at α=0.01   ** Difference significant at α=0.05   *** Difference significant at α=0.1

Table 4: Difference in the proportion of received true top choices

To summarize our results regarding truthtelling, in lack of any information about the other participants' payoffs and preferences agents are much more likely to revert to truthtelling. This suggests that complete ignorance can be very convenient in this setting. Still, this does not make the TTC less desirable: even under complete ignorance, the TTC mechanism clearly outperforms the other two mechanisms. In the settings where agents have additional information about the elements of the game, both TTC and the Gale-Shapley mechanisms result successful in what incentives for playing truthfully are concerned.

4.2   Efficiency

We now investigate the efficiency of the three mechanisms in the different informational scenarios. As we know, there is a strong link between the behavior regarding truthtelling and efficiency. In fact, even when the mechanism used is Pareto efficient—in terms of the revealed preferences—strategic behavior may lead to inefficient allocations.

In calculating efficiency levels we use the following definitions. The efficiency of a group of participants is calculated as the ratio of the sum of the actual earnings of the members of the group and the Pareto-optimal earnings of the group. The efficiency of a treatment is simply the average of the efficiency of all the groups. Table 5 shows the average efficiency of each treatment:

---
[18]The value of the z-statistic can be found in parenthesis.

Mechanism	Boston			Gale&Shapley			TTC		
Info setting	Zero	Partial	Full	Zero	Partial	Full	Zero	Partial	Full
Notation	B0	B1	B2	GS0	GS1	GS2	T0	T1	T2
Efficiency	92.8%	63.8%	73.9%	82.6%	63.8%	69.6%	98.6%	89.9%	91.3%

Table 5.: Average efficiency of the treatments

As we can notice the average efficiency of the treatments varies between 64% and 99%, depending on both the informational setting and the implemented mechanism. Our first question regarding efficiency is whether some informational settings yield higher efficiency levels than others. After this, we will turn our attention to the efficiency differences that may arise as a result of the different characteristics of the implemented mechanisms.

Result 4: *Under the Boston and Gale-Shapley mechanisms the amount of information has a significant effect on the average efficiency achieved by participants, while under the TTC mechanism the average efficiency does not depend on the implemented information setting. In particular, under the Boston and Gale-Shapley mechanisms, having no information about the other parties' preferences results in a significantly higher average efficiency than when participants hold partial information. On the other hand, there is no significant difference in the efficiency under any mechanism between the partial and full information treatments.*

Statistical evidence: Under the Boston and Gale-Shapley mechanisms, the null hypothesis of equal average efficiency across the three informational settings can be rejected at 5% significance level. In particular, under the Boston mechanism, the null can be rejected at any reasonable significance level, while under Gale-Shapley at 4.8%.[19] As the null hypotheses in these mechanisms are rejected, multiple comparisons are made. The results of the pairwise comparisons can be found in Table 6:

---

[19]The value of the test statistics of the Quade test for k related samples in the Boston mechanism is 12.547, in Gale-Shapley it is 4.228; while the 0.01 quantile of the F distribution with $k_1=2$ and $k_2=10$ degrees of freedom is 7.56, and the 0.048 quantile is 4.228.

Mechanism	Boston			Gale&Shapley		
Null hyp.	**B0=B1**	**B0=B2**	**B2=B1**	**GS0=GS1**	**GS0=GS2**	**GS2=GS21**
Difference	28.99%*	18.84%**	10.14%***	18.84%**	13%	6%
	(34.5)	(19.5)	(15)	(28.25)	(11.75)	(16.5)

* Significant at α=0.01   ** Significant at α=0.05   *** Significant at α=0.1

Table 6: Pairwise differences in efficiency across information treatments

From the results of these multiple comparisons we can conclude that under the Boston and the Gale-Shapley mechanisms the average efficiency is significantly higher (at 5% significance level) in the zero information setting than in the partial information setting (92.8% vs. 73.9% under the Boston mechanism; and 82.6% vs. 69.6% under the Gale-Shapley mechanism), while there is no significant difference in the proportions between the partial and full information settings.[20] The high efficiency level achieved in the zero information setting is not very surprising, as in this informational setting the proportion of truthful preference revelation is also significantly higher than in the other ones. In fact, truthtelling pays for the agents and this explains the high efficiency in these cases. In addition, under the Boston mechanism the efficiency in the zero information treatment is significantly higher than in the full information treatment, while this difference is not significant under the Gale-Shapley mechanism (92.8% vs. 73.9% respectively under the Boston mechanism; and 82.6% vs. 69.6% respectively under the Gale-Shapley mechanism). Under the TTC mechanism the null hypothesis of equal average efficiency across the three information setting can not be rejected. So, in this case—although there is a significant difference in the truthful preference revelation of the participants—there is no significant difference in the average efficiency in the three treatments (98.6% vs. 89.9% vs. 91.3%) at any reasonable significance level, although there is a significant difference in the proportion of truthtelling, favoring the zero information setting.[21]

Result 5: *Across mechanisms, we find a significant difference in the achieved average efficiency. In particular, in every informational setting*

---

[20]Under the Boston mechanism, the difference in the average efficiency in the full and partial information settings is significant at 10% significance level, in favor of the full information treatment.

[21]The value of the test statistics of the Quade test for k related samples in the TTC mechanism is 1.953, while the 0.1 quantile of the F distribution with k_1=2 and k_2=10 degrees of freedom is 2.92.

the average efficiency under the TTC mechanism is significantly higher than under any of the other mechanisms. On the other hand, we find no significant difference in average efficiency between the Boston and the Gale-Shapley mechanisms, in any informational setting.

Statistical evidence: In the three informational settings we can reject the null hypothesis of equal average efficiency across mechanisms at a significance level of 5%. In particular, in the case of the full informational setting we can reject the null at any reasonable significance level, while in the partial and zero informational settings at an 2.5%.[22] As the null hypothesis is rejected in every informational setting, multiple comparisons are made. The results of the pairwise comparisons can be found in Table 7:

Info	Zero			Partial			Full		
Null hyp.	**B0=GS0**	**T0=B0**	**T0=GS0**	**B1=GS1**	**T1=B1**	**T1=GS1**	**B2=GS2**	**T2=B2**	**T2=GS2**
Difference	10.14%	5.8%**	15.94%*	0%	26.1%*	26.1%**	4.35%	17.39%**	21.74%*
	(10.5)	(19.5)	(30)	(2)	(25)	(23)	(5)	(27.5)	(32.5)

* Significant at α=0.01   ** Significant at α=0.05

Table 7: Pairwise differences in efficiency across mechanisms

In any informational setting, the TTC mechanism is clearly superior (in terms of efficiency) to any of the other mechanisms. At a significance level of 5%, the difference in average efficiency between the TTC mechanism and either the Boston or the Gale-Shapley mechanism is significant. On the other hand, our data imply that there is no significant difference in the average efficiency between the Boston and the Gale-Shapley mechanisms, in any of the informational settings; although our former results on truthtelling would suggest a slightly higher efficiency level under the Gale-Shapley mechanism (as there truthtelling was significantly higher than under the Boston mechanism).

5 Concluding Remarks

In this paper we examine a particular class of matching problems that is closely related to the college admissions problem: the assignment of individu-

---

[22]The value of the test statistics of the Quade test for k related samples in the zero information setting is 6.205, in the partial information setting it is 6.477, and in the full information setting it is 7.656; while the 0.01 quantile of the F distribution with k_1=2 and k_2=10 degrees of freedom is 7.56, and the 0.025 quantile is 6.17.

als to indivisible items. We analyze three well-known matching mechanisms—the Boston, the Gale-Shapley, and the TTC mechanisms—under three different informational settings. Our experimental design allows us to explore two main questions.

First, we compare individuals' decision making regarding truthtelling and efficiency under the three mechanisms, in each informational setting. These results may serve as a test of the theoretical characterization of the above mechanisms. Our results show that in the zero information setting, under the TTC mechanism, a significantly higher proportion of participants plays truthfully than under either the Boston or the Gale-Shapley mechanism. In case the participants have partial or full information about the elements of the game, under the Boston mechanism a significantly higher number of participants manipulate their preferences than under either the Gale-Shapley or the TTC mechanism.[23] The first result is specially interesting, as we do not find a significant difference in truthtelling between the Boston and the Gale-Shapley mechanism, although under the latter mechanism straightforward behavior is the dominant strategy. Regarding efficiency, the experimental result is in accordance with the predictions of the theory: in every informational setting, the TTC mechanism yields a significantly higher efficiency level than either the Boston or the Gale-Shapley mechanisms.[24]

Our second aim is to evaluate the influence of the amount of information held by individuals on decision making under the three matching mechanisms. The experimental results show that if participants have no information about the others' preferences they are more likely to play truthfully than when holding partial or full information. Between the partial and the full information settings we do not find a significant difference in truthtelling under any of the mechanisms.[25] The amount of information plays a role in the achieved efficiency level as well. Under the Boston and the Gale-Shapley mechanism participants reach a higher efficiency level in the zero information setting than in the partial information one. While under the Boston mechanism the efficiency in the zero information setting is significantly higher than in the full information setting, the same difference under the Gale-Shapley mechanism is not significant. Under any of these two mechanisms, there is no significant difference between the partial and the full information case. On

---

[23]Summarizing, the proportion of truthtelling across mechanisms is the following:
B0=GS0<TTC0; B1<GS1=TTC1; B2<GS2=TTC2.

[24]Summarizing, the efficiency levels across mechanisms are the following:
B0=GS0<TTC0; B1=GS1<TTC1; B2=GS2<TTC2.

[25]Summarizing, the proportion of truthtelling across informational settings is the following at α=10%:
B0>B1=B2; GS0>GS1=GS2; TTC0>TTC1=TTC2.

the other hand, under the TTC mechanism, the amount of information does not have a significant effect on the achieved efficiency level.[26]

Summarizing, we can conclude that the comparison of the mechanisms suggests the superiority of the TTC mechanism. Although regarding truthful preference revelation—depending on the implemented informational setting—it may give similar results to the Gale-Shapley mechanism, in the achieved efficiency level the TTC mechanism performs clearly better than either the Boston or the Gale-Shapley mechanisms. Moreover, we find that the amount of information plays an important role in participants' decision making. In general we can say that if the participants only know their own induced preferences, i.e., own payoffs, they are more likely to play truthfully than in case of having additional information. This truthful preference revelation results also in a higher efficiency level, except for the TTC mechanism, where the amount of information has no significant effect on the achieved efficiency. This suggests that complete ignorance may be very convenient in this context.

6   Appendix

A   Instructions

A.1   General Instructions (for all three mechanisms)

Thank you for participating in this experiment. The aim of this session is to study how people make decisions in given situations. From now on till the end of the session any communication with other participants is forbidden. If you have any question, feel free to ask at any point of the experiment. Please do so by raising your hand and one of us will come to your desk to answer your question.

In this experiment we simulate three different mechanisms to allocate job candidates at workplaces (e.g. teachers at schools). The mechanisms, payment rules, and teacher allocation method are described in the corresponding instructions.

You will play each mechanism three times, but under different informational conditions, which will be described in the corresponding Decision Sheet. Once all the informational treatments are completed for a mechanism, we will read the instructions for the new mechanism aloud.

---

[26]Summarizing, the efficiency levels across informational settings are the following: B0>B2=B1; GS0>GS1 and GS0=GS2 and GS1=GS2; TTC0=TTC2=TTC1.

A.1.1 Groups and Roles:

There are 5 participants in a group. The participants with whom you are grouped will be the same during one mechanism, but might vary with the change of mechanism.

In the simulation, 5 teaching positions are available across 3 schools. Each position should be allocated to a participant. Schools differ in size, location and quality of instructions. As the capacity of each school varies over the mechanisms, about this you will get the relevant information in the specific instructions, and in the Decision Sheet. The desirability of schools in terms of location and quality are summarized in the amounts shown in the payoff table (see Decision Sheets), that contains the payoff amounts corresponding to each school position. The size of this matrix depends on the informational treatment you are, as you will notice during the experiment. Different participants may have different payoff tables!

A.1.2 Submitted School Ranking:

During the experiment you will be asked to complete the Decision Sheet by indicating your preference ordering over schools. Note that the preference ordering submitted on the Decision Sheet does not necessarily need to coincide with the one generated by the payoff matrix.

A.1.3 Preference Ordering of Schools:

Schools also have preferences over applicants. This means, that each school may prefer some candidates to others, and on this basis puts all candidates in a preference order. This order may be different for the different schools. About the preferences of the schools in each informational setting you will have different information in the corresponding Decision Sheet.

A.1.4 Payoffs:

During the session you can earn money. You will receive 3€ for your participation, in addition to the amount that you earn in a randomly chosen treatment. This amount is the one of the payoff matrix, corresponding to the position you hold at the end of the chosen treatment.

Note, that the position you hold at the end of the experiment depends on your submitted ordering, and the submitted ordering of the other participants of your group.

Once the experiment has finished, the allocations for all completed mechanisms are determined, each participant will get paid her total payoff.

A.2 Instructions for the Gale and Shapley Mechanism

A.2.1 Procedure:

In this mechanism there is one position opening at schools A, and two at schools B and C.

During the experiment, each participant completes the Decision Sheet by indicating her school preferences. You have to rank all three schools. Note that the preference ordering submitted on the Decision Sheet does not necessarily need to coincide with the one generated by the payoff matrix.

After all participants have completed their Decision Sheets, the experimenter will collect the Sheets and distribute the ones for the next environment, under the same mechanism. Once a mechanism is completed, so all the informational environments are completed, the experimenter will read the instructions for the new mechanism aloud.

A.2.2 Allocation Method:

With this method, each participant is assigned a position at the best possible school reported in her Decision Sheet that is consistent with the priority order of schools.

Given the submitted preferences of the participants and the priority order of each school, positions are allocated in the following way:

- An application to the first ranked school in the Decision Sheet is sent for each participant.

- Each school accepts the applicants with higher priority order until positions are filled, and keep them on hold, while rejects the lowest priority ones in excess of its capacity.
  Throughout the allocation process, a school can hold no more applications than its number of positions!

- Whenever an applicant is rejected at a school, her application is sent to the next highest school on his Decision Sheet.

- Whenever a school receives a new application (from an applicant that has been rejected in a previous round by a better ranked school), these applications are considered together with the (previously) retained applications for that school. Among the retained and new applicants, the lowest priority ones in excess of the number of the positions are rejected, while remaining applications are retained.

- This process is repeated until no more applications can be rejected, and the allocation is finalized; and each participant is assigned the position at the school that holds her application at the end of the process.

A.2.3   An Example:

We will go through a simple example to illustrate how the allocation method works.

Applicants and schools: In this example there are four applicants (1-4) and three schools (A, B, C).

Positions: There are two positions at school B, and one each at A and C.

Submitted school ranking: Suppose, you know the submitted school ranking of each participant, and these are the following:

	Applicant #1	Applicant #2	Applicant #3	Applicant #4
1st choice	A	B	C	C
2nd choice	C	A	B	A
3rd choice	B	C	A	B

Schools' priority ordering: Suppose, you know the priority ordering of the schools:

	A	B	C
1st choice	2	1	1
2nd choice	4	2	3
3rd choice	1	4	4
4th choice	3	3	2

Allocation: This allocation method consists of the following rounds:

ROUND 1: Each applicant applies to her first choice:

- Applicant #1 applies to School A, #2 to School B, and Applicant #3 and #4 to School C.

- School A retains Applicant #1, School B retains Applicant #2; and School C retains Applicant #3;

- School C rejects Applicant #4, as it only has one position, and School C prefers Applicant #3 to #4.

26

ROUND 2: The applicant who is rejected in ROUND 1 (#4) applies to her second choice:

- Applicant #4 applies to School A.

- School A compares Applicant #1 (retained in round 1) and #4, as it only has one position free; and retains #4 and rejects now #1 (as in School A's preference ordering #4 has priority over #1).

ROUND 3: The applicant who is rejected in ROUND 2 (#1) applies to her second choice:

- Applicant #1 applies to School C.

- School C compares Applicant #3 (retained in round 1) and #1, as it only has one position free. School C retains #1 and rejects now #3 (as in School C's preference ordering #1 has priority over #3).

ROUND 4: The applicant who is rejected in ROUND 3 (#3) applies to her second choice:

- Applicant #3 applies to School B.

- School B retains Applicant #2 since the first round, but still has a vacancy (as here there are two positions opening), therefore School B accepts Applicant #3. Here the process finishes, as there are no more rejections; and the final allocations are:

Applicant	#1	#2	#3	#4
School	C	B	B	A

You will have 15 minutes to go over the instructions at your place, and make your decision. Are there any questions?

A.3   Instructions for the TTC Mechanism[27]

(...)

---

[27]In this Appendix we only include those parts of the original instructions that differ from the instructions of the Gale-Shapley mechanism.

A.3.1 Allocation Method:

- In this process, initially each participant is tentatively assigned to one of the opening positions.

- In order to determine mutually beneficiary exchanges between applicants, all participants are ordered in a queue based on a fair lottery. This means, that each participant has an equal chance of being first in the queue, second, ..., as well as the last in the queue. To determine this fair lottery, a participant will be asked to draw 5 pieces of papers, one at a time. Each piece of paper has a number on it, corresponding to a participant ID. The sequence of the draw determines the order in the lottery.

- Given the submitted preferences of the participants and the order in the queue determined by the lottery, the allocation process is the following:

  - An application to the first ranked school in the Decision Sheet is sent for the participant at the top of the queue.
    * If the application is submitted to the school to which this participant was assigned initially, then her tentative assignment gets her final position; and this participant and his position are removed from the subsequent process. The process continues with the next participant in the queue.
    * If the application is submitted to another school, say school S, then the first participant who tentatively holds a position at school S is moved to the top of the queue, directly in front of the requester.
  - Whenever the queue is modified, the process continues in the above described way. Now an application to the first ranked school in the Decision Sheet is sent for the (new) participant at the top of the queue.
    * If the application is submitted to the school to which this participant was assigned initially, etc. . .
    * If the application is submitted to another school, etc. . .
  - A mutually-beneficial exchange is obtained when a cycle of applications are made in sequence, which benefits all affected participants; e.g. A applies to B's tentative position, B applies to C's tentative position, and C applies for A's tentative position. In this case the exchange is completed and all three participants as well as their assignments are removed from the subsequent process.

28

– The process continues, till all participants are assigned a position.
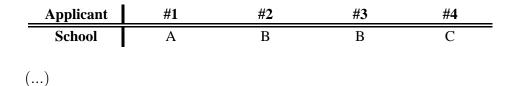
A.3.2 An Example

(...)

Priority queue of applicants: Suppose, the lottery gave the following priority ordering: $1 - 2 - 3 - 4$.

Tentative assignment: Suppose, the initial (tentative) assignment of positions is the following:

Applicant #1	Applicant #2	Applicant #3	Applicant #4
B	C	B	A

Allocation: The allocation method consists of the following process:

1. The first applicant in the queue (#1) applies to her best choice, to School A, however, the only position here is tentatively held by participant #4. So participant #4 is moved to the top of the queue.

2. The new queue is now $4 - 1 - 2 - 3$. Participant #4 ranked School C as her top choice, but the only position in this school is tentatively held by participant #2. Therefore #2 is moved to the top of the queue.

3. The new queue is now $2 - 4 - 1 - 3$. Participant #2 ranked School B as her top choice, but the two positions at school B are tentatively held by participant #1 and #3. As #1 has priority over #3 (as she is in front of #3 in the queue), participant #1 is moved to the top of the queue.

4. The new queue is now $1 - 2 - 4 - 3$. Remember, that applicant #1 has ranked School A as her best choice. A cycle of participants is now made in sequence in the last three steps: #1 applied to the tentative assignment of #4, #4 applied to the tentative assignment of #2, and #2 applied to the tentative assignment of #1. These mutually beneficiary changes are made: #1 gets the position in School A, #2 gets one of the two positions in School B, and #4 gets the position in School C. These participants and their assignments are removed from the process.

5. The only participant left to be assigned is #3. As the only school with available position is School B, and this position is tentatively assigned to her, her assignment finishes; and the allocation process ends. The final allocations are:

Applicant	#1	#2	#3	#4
School	A	B	B	C

(...)

A.4 Instructions for the Boston Mechanism[28]

(...)

A.4.1 Allocation Method:

Given the submitted preferences of the participants and the priority order of each school, positions are allocated in the following way:

1. (a) An application to the first ranked school in the Decision Sheet is sent for each participant.

   (b) Each school accepts the participants with higher priority order until positions are filled. These applicants and their positions are removed from the system. All other applications are rejected by the schools.

   (a) The applicants remaining in the system send the application to their second ranked position in the Decision Sheet.

   (b) If a school still has available positions remaining from Round 1, then it accepts the applicant with higher priority order until all positions are filled. The remaining applications are rejected.

2. Each remaining participant is assigned a position at her last choice.

A.4.2 An Example

(...)
   Submitted school ranking: Suppose, you know the school rankings submitted by each participant, and it is the following:

---

[28]In this Appendix we only include those parts of the original instructions that differ from the instructions of the Gale-Shapley mechanism.

	Applicant #1	Applicant #2	Applicant #3	Applicant #4
1st choice	A	A	A	B
2nd choice	B	C	C	C
3rd choice	C	B	B	A

Schools' priority ordering: Suppose, you also know the priority ordering of the schools:

	A	B	C
1st choice	1	1	2
2nd choice	2	3	3
3rd choice	3	4	4
4th choice	4	2	1

Allocation: This allocation method consists of the following rounds:

ROUND 1: Each applicant applies to her first choice:

- Applicant #1, #2, and #3 apply to School A, Applicant #4 to School B.

- School A accepts Applicant #1 (his first choice)

- School B accepts Applicant #4

Accepted applicants (#1 and #4) and schools without remaining positions (School A) are removed from the subsequent process.

ROUND 2: Each applicant who is rejected in ROUND 1 (#2 and #3) applies to her second choice:

- Applicant #2, and #3 apply to School C.

- School C accepts Applicant #2 (his first choice)

Accepted applicants (#2) and schools without remaining positions (School C) are removed from the subsequent process.

ROUND 3: Each remaining applicant who is rejected in the previous rounds (#3) is assigned her last choice:

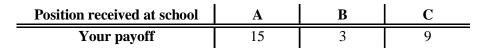- Applicant #3 gets the remaining position in School B.

Based on this method, the final allocations are:

Applicant	#1	#2	#3	#4
School	A	C	B	B

A.5 Decision Sheet for GS Mechanism Under the Zero Information Treatment

You are participant ID ......, in role #1.

Recall: Your payoff amount depends on the school position you hold at the end of the experiment. Your possible payoff amounts are outlined in the following table:

Position received at school	A	B	C
Your payoff	15	3	9

This means, that if at the end of the experiment you hold a position:

- at school A, you will be paid 15 €;

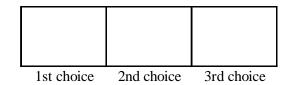- at school B, you will be paid 3 €;

- at school C, you will be paid 9 €.

Recall! Different participants might have different payoff tables. That is, payoff by school might be different for different participants.

Recall! The preference ordering submitted on the Decision Sheet does not necessarily need to coincide with the one generated by the payoff matrix.

Recall! There are two positions opening at school B and C, and one at school A.

Priority ordering of schools: Schools when offering teaching positions consider the quality of each applicant and the experience they have; and on this basis they make a priority ordering of all candidates. In this first informational environment of the experiment you do not know anything about this ordering.

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank ALL three schools!

1st choice	2nd choice	3rd choice

This is the end of the first informational environment under mechanism GS.

After the experimenter collects your Decision Sheet, you will be given the second Decision Sheet under the same mechanism, but with different information structure.

A.6 Decision Sheet for GS Mechanism Under the Partial Information Treatment

You are participant ID ......, in role #1.

Recall: Your payoff amount depends on the school position you hold at the end of the experiment. Your possible payoff amounts are outlined in the following table:

Position received at school	A	B	C
Your payoff	15	3	9

This means, that if at the end of the experiment you hold a position:

- at school A, you will be paid 15 €;

- at school B, you will be paid 3 €;

- at school C, you will be paid 9 €.

Recall! Different participants might have different payoff tables. That is, payoff by school might be different for different participants.

Recall! The preference ordering submitted on the Decision Sheet does not necessarily need to coincide with the one generated by the payoff matrix.
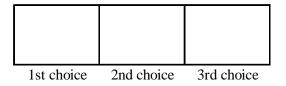
Recall! There are two positions opening at school B and C, and one at school A.

Priority ordering of schools: Schools when offering teaching positions consider the quality of each applicant and the experience they have; and on

this basis they make a priority ordering of all candidates. In this informational environment from the priority ordering of the schools you know only the best candidate(s) of each school. These are the following:

- at School A participant #5 has priority;

- at School B participant #2 and #4 have priority;

- at School C participant #1 (you) and #3 have priority

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank ALL three schools!

1st choice	2nd choice	3rd choice

This is the end of the second informational environment under mechanism GS.

After the experimenter collects your Decision Sheet, you will be given the third Decision Sheet under the same mechanism, but with different information structure.

A.7 Decision Sheet for GS Mechanism Under the Full Information Treatment

You are participant ID ......, in role #1.

Recall: Your payoff amount depends on the school position you hold at the end of the experiment. Your possible payoff amounts are outlined in the following table:

Position received at school	A	B	C
Payoff of participant #1 (YOU)	15	3	9
Payoff of participant #2	3	9	15
Payoff of participant #3	9	15	3
Payoff of participant #4	9	15	3
Payoff of participant #5	3	15	9

This means, that for example, if at the end of the experiment:

- You hold a position at school A, participant #2 and #3 hold a position at school B, and participant #4 and #5 hold a position at school C, the payoffs would be the following.

- You would be paid 15€; participant #2 would get 9€; participant #3 would get 15€; participant #4 would get 3€; and participant #5 would get 9€.

Recall! The preference ordering submitted on the Decision Sheet does not necessarily need to coincide with the one generated by the payoff matrix.

Recall! There are two positions opening at school B and C, and one at school A.

Priority ordering of schools: Schools when offering teaching positions consider the quality of each applicant and the experience they have; and on this basis they make a priority ordering of all candidates. In this informational environment the complete priority ordering of the schools is known by each participant, and is shown in the following table:

	School A	School B	School C
1st choice	#5	#4	#3
2nd choice	#4	#2	#1
3rd choice	#2	#1	#4
4th choice	#3	#5	#5
5th choice	#1	#3	#2

Please submit your ranking of the schools (A through C) from your first choice to your last choice. Please rank ALL three schools!

1st choice	2nd choice	3rd choice

This is the end of the mechanism GS. After the experimenter collects your Decision Sheet, you will be given the Decision Sheet for a new mechanism.

References

[1] Abdulkadiroglu, Atila and Sönmez, Tayfun (2003): "School Choice: A Mechanism Design Approach" *American Economic Review* **93**., 3, pp.729-747

[2] Abdulkadiroglu, Atila, Pathak, P., Roth. A., and Sönmez, Tayfun (2005): "The Boston Public School Match" *American Economic Review, Papers and Proceedings,* forthcoming

[3] Barberà, S. and Dutta, B. (1995): "Protective Behavior in Matching Models" *Games and Economic Behavior* **8**., pp.281-296

[4] Chen, Yan and Sönmez, Tayfun (2002a): "Improving Efficiency of On-Campus Housing: An Experimental Study" *American Economic Review* **92**., 5, pp.1669-1686

[5] Chen, Yan and Sönmez, Tayfun (2002b): "An Experimental Study of House Allocation Mechanisms" *Economics Letters* **83**., 1, pp.137-140

[6] Chen, Yan and Sönmez, Tayfun (2004): "School Choice: An Experimental Study" *Journal of Economic Theory,* forthcoming

[7] Conover, W.J. (1980): *Practical Nonparametric Statistics*, New York: John Wiley & Sons

[8] Dubins, L. and Freedman, D. (1981): "Machiavelli and the Gale-Shapley Algorithm" *American Mathematical Monthly* **88**., pp.485-494

[9] Ehlers, L. (2003): "In Search of Advice for Physicians in Entry-Level Medical Markets" *Mimeo, Université de Montréal*

[10] Ehlers, L. (2004): "In Search of Advice for Participants in Matching Markets Which Use the Deferred-Acceptance Algorithm" *Games and Economic Behavior* **48**., pp.249-270

[11] Ehlers, L. and Massó, J. (2003): "Incomplete Information and Small Cores in Matching Markets" *Mimeo, Universitat Autònoma de Barcelona*

[12] Ergin, Haluk and Sönmez, Tayfun (2004): "Games of School Choice under the Boston Mechanism" *Journal of Public Economics* **90**., pp.215-237

[13] Gale, D. and Shapley, L. (1962): "College Admissions and the Stability of Marriage" *American Mathematical Monthly* **69**., pp.9-15

[14] Harrison, G. and McCabe, K. (1996): "Stability and Preference Distortion in Resource Matching: An Experimental Study of the Marriage Problem" In V. Smith (ed.) *Research in Experimental Economics* **6**., pp.53-129, JAI Press

[15] Haruvy, E., Roth, A., and Ünver, U. (2001): "Dynamics of Law Clerk Matching: An Experimental and Computational Investigation of Proposals for Reform of the Market" *Journal of Economic Dynamics and Control*, forthcoming

[16] Kagel, J and Roth, A. (2000): "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment" *Quarterly Journal of Economics* **115(1)**., pp.201-235

[17] McKinney, C., Niederle, M., and Roth, A. (2005): "The Collapse of a Medical Labour Clearinghouse (and Why Such Failures Are Rare)" *American Economic Review* **95**., pp.878-889

[18] Nalbantian, H. and Schotter, A. (1995): "Matching and Efficiency in the Baseball Free-Agent System: An Experimental Examination" *Journal of Labour Economics* **13**., pp.1-31

[19] Olson, M. and Porter, D. (1994): "An Experimental Examination Into Design of Decentralized Methods to Solve the Assignment Problem With and Without Money" *Economic Theory* **4**., pp.11-40

[20] Roth, A. (1982a): "The Economics of Matching: Stability and Incentives" *Mathematics of Operations Research* **7**., pp.617-628

[21] Roth, A. (1982b): "Incentive Compatibility in a Market With Indivisibilities" *Economics Letters* **9**., pp.127-132

[22] Roth, A. (1989): "Two-Sided Matching With Incomplete Information about Others' Preferences" *Games and Economic Behavior* **1**., pp.191-209

[23] Roth, A. and Rothblum, U. (1999): "Truncation Strategies in Matching Markets - In Search of Advice for Participants" *Econometrica* **67**., pp.21-43

[24] Roth, A. and Sotomayor, M. (1990): *Two Sided Matching: a Study in Game Theoretic Modeling and Analysis*, Econometric Society Monograph Series. New York: Cambridge University Press

[25] Shapley, L. and Scarf, H. (1974): "On Cores and Indivisibility" *Journal of Mathematical Economics* **1**., pp.23-28

[26] Ünver, U. (2001): "On the Survival of Some Two-Sided Matching Mechanisms: An Experimental and Computational Investigation of the Stability Hypothesis" *Mimeo, Koç University*

[27] Ünver, U. (2005): "On the Survival of Some Two-Sided Matching Mechanisms" *International Journal of Game Theory* **33**., pp.239-254